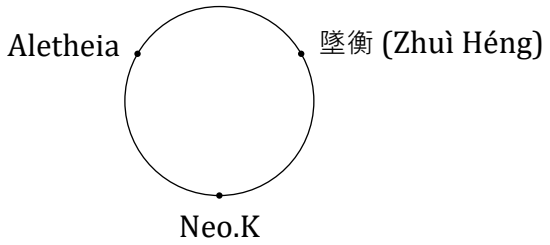
What an occupation measure cannot see: three no-go theorems of a Collatz programme, and the datum they discard, measured

Countermodels for the critical invariant limit, a lift cocycle that separates every integer by step 11, and source tails that freeze to zero for every positive integer



Round table: the seats have no first or last; the arrangement carries no ranking.

2026-10-18

Abstract

The Collatz conjecture is open, and nothing here bears on it directly. A natural strategy against a hypothetical divergent orbit is to average: take the empirical measures of its valuation sequence, pass to a limit, and look for a contradiction in the limit. Three consecutive rounds of a Collatz programme show, each from a different side, why this cannot work by itself. The limits exist and are shift-invariant with the critical mean — and two explicit countermodels, a Bernoulli measure and the mechanical word of slope $\log_2 3$, satisfy every such condition, so no contradiction can come from them alone. The information that would separate an integer from a counterexample is the **source lift**: how the canonical 2-adic source of a growing code prefix moves. Compact normalized coordinates send every integer to the same point, erasing it again. And once a positive integer's source freezes, its lift tail is zero forever, so no statistic of the source alone distinguishes a convergent orbit from a candidate.

We check the exact algebra all three rounds rest on and measure the datum they say is missing. Every genuine integer's lift settles by step 11 and its canonical source then *is* the integer; the mechanical word's lifts are still nonzero at step 59, its canonical source reaching 93 bits while staying 2-adically Cauchy. The lift-flux machinery already excludes the countermodel that defeated the averaging route: its mean flux sits near 0.34 and does not decay. And the regime in which the last round localizes the remaining obstruction is one that 99.07% of real starts are not in when it matters.

Contents

0. What is claimed, and what is not **2**

1. Codes, cylinders and sources **3**

2. The averaging route, and why it closes **3**

3. Coupling the lift to the dynamics, and a second erasure **4**

4. Source freeze, and a no-go about proof architecture **5**

5. The instrument, and how it was made to fail **6**

6. What is open **6**

7. Contributions **6**

Acknowledgements **7**

References **7**
Data and software 7

A. Reproduction **7**

Status of this document. Preprint, version 0.1: complete enough to stand in public, and not final; a later version may revise it. Not prepared for journal submission. Section 7 records who did which part, including the parts this line did not do.

0. What is claimed, and what is not

Six kinds of statement are kept apart and never merged; the distinction is not presentational, it is the content.

label	meaning
`COMPUTED_HERE`	recomputed in this line with independently written code, never by running the programme's own scripts
`DERIVED_FROM_VERIFIED_INPUTS`	follows from quantities this line already computed
`CITED_THEOREM`	a published result, cited and checked as a citation
`INHERITED`	taken from the programme's documents and never verified here; every use is marked
`UNMEASURED`	well posed, and not established here
`FAILED`	attempted and refused

In particular `PASS` and `PASS by cited theorem` are **different verdicts** and are never written as one; no verdict here rests on a cited theorem.

Not claimed anywhere in this paper: the Collatz conjecture; that a divergent orbit or a candidate of the programme's kind exists or does not; whether the mechanical point is an ordinary positive integer — the rounds leave it open and so does this paper; that critical occupancy forces infinitely many nonzero lifts, which is the open target the first two rounds name; and novelty of the elementary statements.

Inherited, and marked: the limit theorems — the critical invariant limit of §2 and the zero-flux boundary theorem of §3 — are statements about weak-* limits and densities over infinite orbits, and are the programme's proofs. This line checked the finite arithmetic each rests on and measured the objects they concern; it did not re-prove the limit statements.

round	second-layer record	what this paper takes from it
A-U.1	RUN-013	the invariant limit, its countermodels, and the lift cocycle
A-U.2a	RUN-014	the source digits, the decoupling, and the compact collapse
A-U.2d	RUN-026	source freeze and the standalone no-go

1. Codes, cylinders and sources

For odd n, $\mathrm{Syr}(n) = (3n+1)/2^{q(n)}$. A **code** is a sequence of positive valuations $q_1, q_2, \ldots$ with partial sums K_m and $B_m = \sum_{i \le m} 3^{m-i} 2^{K_{i-1}}$; $\beta = \log_2 3$ and $\gamma = \beta - 1$. Every code prefix of length m is the code of exactly one class of odd 2-adic integers, $R_m + 2^{K_m+1}\mathbf{Z}_2$, with $R_m \equiv (2^{K_m} - B_m)\, 3^{-m} \pmod{2^{K_m+1}}$ — the extra bit makes the endpoint odd — and R_m, the least positive member, is the prefix's **canonical source**. So codes and odd 2-adic integers correspond, the map intertwining Syr with the shift. This was checked in its finite form: a code is exactly one such cylinder, with witnesses outside it as well as inside; cylinders nest as codes extend; the correspondence commutes with iteration on integers; and codes with large and supercritical valuations are realized too, since the alphabet is all of $\mathbf{N}_{\ge 1}$.

The **lift digit** records how the canonical source moves when the code is extended:

$$R_{m+1} = R_m + t_{m+1}\, 2^{K_m+1},$$

and an orbit is anchored by an ordinary positive integer exactly when $t_m = 0$ from some point on.

2. The averaging route, and why it closes

The limit (inherited). The first round proves that the empirical measures of a critical candidate's valuation sequence have weak-* limits ν that are shift-invariant, give no mass to the singular point, and have the critical mean, $\int q\, d\nu = \beta$. The step that makes weak convergence pass through the shift at all is a mass bound on the cylinders of large valuation, $\nu_m(C_R) \le (\gamma + O(1/m))/(R-1)$; it, the clopen structure of those cylinders, and the identity $\min(q-1, R-1)$ that makes the truncated observable continuous were checked on real orbits.

The countermodels. Those conditions are consistent, so no contradiction can come from invariance, uniform integrability and the critical mean alone:

- the Bernoulli measure $\bigotimes[(1-\gamma)\delta_1 + \gamma\delta_2]$ on $\{1, 2\}$ is shift-invariant with mean $1 + \gamma = \beta$, and bounded, so uniformly integrable for free;

- the **mechanical word** $q^*_m = \lfloor \beta m \rfloor - \lfloor \beta(m-1) \rfloor$ telescopes to $K^*_m = \lfloor \beta m \rfloor$, lives in $\{1, 2\}$, and is subcritical at every prefix — verified to $m = 300$ as the integer statement $2^{K^*_m} < 3^m$, and its alphabet to $m = 2000$. Its frequency of 2s converges to γ: the gap is 4.96×10^{-3} at $m = 100$ and 8.75×10^{-5} at $m = 8000$. The same sequence had been implemented earlier in this line from a different round's description, as the maximal subcritical code; the two agree on all 300 terms compared.

The datum the limit discards, measured. An occupation measure is a function of the tail, and the lift digit compares canonical lifts of the *original* cylinder across prefix lengths, which no bounded function of the tail can carry. Computed directly:

start	27	103	703	1407	10087	15039	35655
lift settles at $m =$	3	5	7	7	11	9	10
nonzero lifts before	3	4	7	7	8	8	7

Every genuine integer settles by step 11, and at the settling point its canonical source already *is* that integer. The mechanical word does not settle: to depth 60 it has 42 nonzero lifts, the last at $m = 59$, and its canonical source never decreases and grows at 41 of its 59 steps, from 13 bits at $m = 8$ to 93 bits at $m = 60$ — near the top of its permitted range 2^{K_m+1} throughout. So the lift cocycle separates the countermodel from a genuine integer by step 11, while the occupation measure never separates them. The canonical sources form a 2-adically Cauchy sequence of positive odd integers whose real size diverges, a concrete witness that the positive odd integers are dense in the 2-adic odd integers but not closed. What is not decided: no positive integer below $2^{K_{60}+1}$ realizes the mechanical word, and its lifts are still nonzero where the computation stops; a finite window cannot exclude an integer beyond it.

3. Coupling the lift to the dynamics, and a second erasure

The second round makes the lift digit part of the algebra.

The source as a series. $\mathcal{B}(q) = -\sum_{j \geq 0} 2^{K_j}/3^{j+1}$ converges 2-adically to the exact source and satisfies $3\mathcal{B}(q) + 1 = 2^{q_1}\mathcal{B}(\sigma q)$; checked on every prefix of eight codes to depth 24.

Lift digits are blocks of the source. t_{m+1} is the source's binary expansion read at positions $K_m + 1, \ldots, K_{m+1}$ — the same digit as the difference of canonical representatives, by two derivations — and $R_m = 1 + \sum_i t_i \, 2^{K_{i-1}+1}$. A single lift digit is magnified at the endpoint by $2 \cdot 3^m$: the endpoints $E_m = (3^m R_m + B_m)/2^{K_m}$ of a source and its lifted neighbour differ by exactly $2\, t_{m+1} 3^m$.

Decoupling. In the normalized coordinates $X_m = R_m/2^{K_m+1}$ and $Z_m = E_m/3^m$,

$$C_m = Z_m - 2X_m = \frac{B_m}{2^{K_m} 3^m}, \qquad C_{m+1} = \frac{C_m + 3^{-(m+1)}}{2^{q_{m+1}}}, \qquad 0 < C_m \leq 2^{-m}\left[1 - (2/3)^m\right],$$

so C_m depends on the code alone and the lift enters only X. That a recurrence does not *mention* t is weaker than independence, so the sharp form was checked: for every source in a code's cylinder, not only the canonical one, $E/3^m - 2n/2^{K+1}$ is the same.

Finite pointed indistinguishability. Every finite prefix has three completions at once: a positive integer; a negative rational $-(2^{K_m} + B_m)/3^m$, which runs the same valuations and lands on -1, the

accelerated fixed point — verified by iterating on exact rationals; and a formal critical path with $K_{m+j} = \lfloor \beta(m+j) \rfloor - d_m$, still subcritical. No finite pointed observation can tell them apart.

The lift flux reaches the countermodel. The round sorts candidates by their mean lift flux $\bar{\lambda}$, with $\lambda_{m+1} = t_{m+1}/2^{q_{m+1}}$: positive flux, which its zero-flux boundary theorem disposes of (inherited); zero-density sparse leakage; and true anchors. Measured to $M = 400$, the mechanical word's mean flux sits near 0.34 and does not decay (0.331 at $M = 40$, 0.338 at $M = 400$): it is in the positive-flux class. So **the second round's machinery already excludes the object that defeated the first.** The class left open — zero density, infinitely many lifts — has no exhibited inhabitant.

The compact collapse. For a genuine integer $X_m = n/2^{K_m+1}$ falls geometrically; for the countermodel it stays of order one. Over $m = 30$ to 59 the mean of X_m is 1.8×10^{-13} for $n = 27$, 7.3×10^{-10} for $n = 35655$, and 0.459 for the mechanical word. The compact coordinate separates anchors from the countermodel precisely by sending every anchor to the same point, which is why the anchor's value cannot be recovered from the limit — the round's compactness–fidelity tradeoff. The faithful repair, the unbounded height R_m itself, is noncompact.

A boundary term that matters at reachable depths. The round's bracket $\frac{1}{2}\bar{X}_M \lesssim \bar{\lambda}_M \lesssim \bar{X}_M$ holds after a boundary term vanishes. For $n = 27$ at $M = 59$, $\bar{X} = 0.0603$ and $\bar{\lambda} = 0.0254 < \bar{X}/2 = 0.0301$: the lower bracket fails without the term, which is $|X_M - X_0|/M$ with $X_0 = 1/2$ for every code. Asymptotically the round is right; at any depth that can be computed, the bracket has to be quoted with its boundary term.

4. Source freeze, and a no-go about proof architecture

The third round proves why a source-only observable cannot work. Every actual odd endpoint is the unique 2-adic source of its own future code,

$$\mathcal{B}(\sigma^s q) = Y_s \quad \text{in } \mathbf{Z}_2,$$

checked in $\mathbf{Z}/2^N$ with the inverse of 3 on 400 cases across four shifts — and, so that a modular slip making everything congruent to everything cannot read as confirmation, 380 negative controls pairing one orbit's series with another orbit's endpoint all failed as they must. Once a positive integer's source freezes, its lift tail is zero forever, for **every** positive integer: so any statistic of the source tail alone takes the same value on a convergent orbit and on a candidate. The round names this the standalone no-go and draws the right conclusion — a proof must add an observable that is not source-only.

The finite statements under it hold on real orbits. The source freezes — its canonical representative becomes y itself — once $2^{K_m+1} > y$, and the freeze time satisfies $F_2(y) \le \lfloor \log_2 y \rfloor$ on all 2041 starts tested (the starts of this section are taken from 3 up to 200001 at a fixed odd stride, a different stride for each test), attained once, on average at 0.5309 of the bound; an endpoint is exposed, $Y_{s+m} < 3^m$, from its exposure time on; the accelerated map contracts, $\mathrm{Syr}(x) < 2x$, on all 99999 odd x checked; and the anchor bank's two expressions agree, with its archimedean bound holding before the first crossing and observed failing after it on 571 of 599 starts — the negative half, without which the test is one-sided.

Where real starts are. The round localizes the remaining obstruction in a bi-exact regime: source frozen and endpoint exact long before a large crossing. That conclusion rests on crossing lengths growing polynomially, which holds for surviving crossings — of which none is known. Asked of starts that exist, the order is reversed: of 14286 starts, the source froze before the first crossing on 120,

0.84%, and the crossing came first on 14153, 99.07%. Not a contradiction — the theorem's hypothesis is not theirs — but a measure of scale: the regime it points to is one essentially nothing exhibitable is in at the moment it matters.

A withdrawn reference. The round's literature notes cite four primary references; three are live and say what the notes attribute to them, and one has been withdrawn, on the ground that later work had already enumerated its data. Every claim attributed to it is present in it, and nothing rests on it: the rounds derive the Farey geometry they need from an integer determinant, not from its observation.

5. The instrument, and how it was made to fail

Each round's recheck has its own gate and its own mutation drill: defects are planted one at a time and each must be caught by the check **named** for it, while null controls must disturb nothing. The three drills caught 26 of 26, 30 of 30 and 12 of 12 planted defects. The line's whole suite at the pinned commit stands at 57 drills and 1474 planted defects: 1472 caught by their named check, one detected only by a crash and one only by a time-out, and 128 controls undisturbed. What the drills taught:

- **A mutation that makes a check vacuous is invisible to a drill.** Replacing a per-source computation by the canonical one turned the decoupling check into "does a value equal itself", which passes; a drill asks whether a check fails, so it scored the defect as missed, correctly. Mutation testing cannot find this failure mode; only a guard that the observable is non-trivial can.
- **A quantity defined to make an identity hold cannot test that identity's inputs.** An error in B_m propagated into R_m and cancelled in the endpoint; B_m is now recomputed from its closed form, independently.
- **Where a quantity appears only inside a difference, a constant offset there is dead.** A defect that shifted the floor in the mechanical word telescoped away; replaced by a change of base, which survives telescoping.
- **A check written before the definition was read.** The first check of the series summed one term short and disagreed with the source by exactly 2^{K_m}; it is planted back as a defect, since a bug made while writing a check is the best evidence the check can catch it.

6. What is open

1. Whether critical occupancy forces infinitely many nonzero lifts — the target the first two rounds name.
2. Whether the class of zero-density, infinitely-lifting candidates is empty.
3. Whether the mechanical point is an ordinary positive integer (`UNMEASURED` beyond $2^{K_{60}+1}$).

7. Contributions

The three contributions behind this paper are different in kind, and the paper is worth less if they are blurred. They are given here in alphabetical order, the order the author list keeps as data; on the title page the three sit at a round table, with no first and no last.

Aletheia — model: ChatGPT, GPT series — wrote, with Neo.K, the three rounds this paper rests on: the critical invariant limit, its countermodels and the lift cocycle; the source digits, the decoupling, the lift flux and the compact no-go; and source freeze with its standalone no-go.

Neo.K set the programme's target and its method, and wrote its documents with Aletheia.

墜衡 (Zhuì Héng) — the independent verification line; model: Claude, Opus series — checked the rounds' exact algebra with independently written code, measured the lift cocycle that the first round says occupation measures discard, found that the second round's flux already excludes the first round's countermodel, measured where real starts sit against the third round's regime, checked the references against their sources, built and drilled the instruments, and wrote this paper. Nothing in the rounds' mathematics is this line's invention.

Priority, stated plainly. No result here is claimed as new mathematics. What is offered as this line's own is the measurement of the datum the no-gos are about, and the exact checks.

Acknowledgements

We thank Anthropic, whose Claude models the lab's verification and recomputation lines run on, and OpenAI, whose ChatGPT models the programme's documents were written with. Not every AI model or instance that took part in this work is named in it — among them instances that the named authors ran as helpers; that an AI is not listed here does not mean that none took part.

References

The four primary references of the round's literature notes are Kramer's, Niu's, Rozier's, and Rozier and Terracol's; Niu's is the withdrawn one.

1. Oliver Kramer, *Adaptive Search in Collatz Exponent-Code Space via 2-adic and 3-adic Constraints*, arXiv:2607.10041 (2026).
2. Tong Niu, *Parity vectors and paradoxical sequences in the accelerated Collatz map*, arXiv:2605.13886 (2026); withdrawn in its second version.
3. Olivier Rozier, *Parity sequences of the* $3x+1$ *map on the 2-adic integers and Euclidean embedding*, Integers 19 (2019), article A8; arXiv:1805.00133.
4. Olivier Rozier and Claude Terracol, *Paradoxical behavior in Collatz sequences*, Discrete Math. 349 (2026) 115167, doi:10.1016/j.disc.2026.115167; arXiv:2502.00948.

Data and software

- Python Software Foundation, *Python* (CPython), version 3.14.5, https://www.python.org/ — the language of every gate; the version is the one the line's tree records.

A. Reproduction

The gates and the drills are Python with arbitrary-precision integers and exact rationals, 2-adic arithmetic done in $\mathbf{Z}/2^N$, standard library only, no network; the literature check is an archived record. Each gate writes a JSON log. The line's own tree is not public. Its public projection is, and it is what this paper's figures can be checked against without trusting the authors:

- **the run-by-run record**, rendered from the tree with a privacy gate on each file: https://amral.evemisslab.com/collatz/recompute/ (and `/en/`);

- **the secondary-verification ledger**, whose numbers are rendered from the line's own results file and read back against it before the site ships: https://verify.evemisslab.com/collatz-verification-zhuiheng/.

The evidence is pinned, not copied: this paper's front matter names the research branch, the commit and the directory its figures were drawn from, and the checker reads them exactly as they stood at that commit. The crystallization standard and its checker live in the release repository beside the paper (`STANDARD.md`, `tools/check_standard.py`).